

UNITED STATES PATENT APPLICATION

FOR

TUNNELING ETHERNET

INVENTORS
SUHAIL NANJI

PREPARED BY:

BLAKELY, SOKOLOFF, TAYLOR & ZAFMAN
12400 WILSHIRE BOULEVARD
SEVENTH FLOOR
LOS ANGELES, CA 90025-1026

(408) 720-8598

04906.P003

TUNNELING ETHERNET

BACKGROUND OF THE INVENTION

Field of the Invention

This invention relates to the area of communication networks. More specifically, the present invention relates to preserving Ethernet frames across a network.

Description of the Related Art

Ethernet headers contain useful data. For example, the source address can be used for network analysis, tracing, billing and accounting purposes. Ethernet frames can be tunneled with the Bridge Encapsulation Protocol (BCP) and the Generic Routing Protocol (GRE).

Figure 1 (Prior Art) is a diagram of transmitting an Ethernet frame with Bridge Control Protocol (BCP) encapsulation. The subscriber's machine 101 (e.g. personal computer (PC)) runs the Internet Protocol (IP) and Ethernet. Data packets are prepared for transmittal to the Internet by prepending IP headers. The subscriber's machine 101 also prepends Ethernet headers onto the data packets to be transmitted. The encapsulation of the data is represented by the protocol stack 103. An encapsulated data packet is transmitted to the customer premise equipment (CPE) 105, such as a cable modem or digital subscriber line (DSL) modem. To preserve the Ethernet frame with BCP, the CPE 105 must run BCP and the Point to Point Protocol (PPP). The CPE 105 encapsulates the data from the subscriber machine 103 with BCP and then PPP. The CPE further prepares the data for transmission by encapsulating the data in delivery protocols (such as 1483 or 1490 bridged circuit over Asynchronous Transmission Mode (ATM), Frame Relay, etc) as indicated by the protocol stack 107. The protocol stack 107 is transmitted to a network element 109. The network element may support BCP and PPP to decapsulate the data packet. The network element 109 is configured as a Layer 2

Tunneling Protocol (L2TP) Access Concentrator (LAC) in order to tunnel the stack or data packet. The LAC 109 encapsulates the data packet with L2TP as indicated by the stack 111 for transmission to an L2TP Network Server (LNS) 113. The LNS 113 must support PPP and BCP to decapsulate the data packet.

Figure 2 (Prior Art) is a diagram of transmitting traffic from multiple subscribers who use different protocols with multiple Generic Routing Encapsulation (GRE) tunnels. Two individual subscribers 201, 202 connect to two separate CPEs 213, 206. The traffic transmitted from the subscribers 201, 202 to the CPEs 213, 206 are IP encapsulated with Ethernet as indicated by the protocol stacks 205, 204. The CPEs 213, 206 encapsulate the subscriber traffic in a delivery protocol, as indicated by the protocol stacks 209 and 210, and transmit the traffic to a network element 217. Another subscriber 203 transmits traffic to a CPE 215 encapsulated with Point to Point protocol as indicated by the protocol stack 207 showing IP over PPP. The CPE 215 encapsulates the subscriber traffic in a delivery protocol as indicated by the protocol stack 211 and transmits the traffic to a network element 217. The network element 217 uses GRE to tunnel the traffic from subscribers 201, 202, and 203 to another network element 227. The network element 217 establishes 2 different tunnels 223, 225 with the network element 227. One tunnel 223 carries both subscribers' Ethernet traffic. The Ethernet traffic is indicated by the protocol stack 221. The other tunnel 225 carries the PPP subscriber traffic indicated by the protocol stack 219. Both GRE tunnels must be transmitted by IP media also indicated by stacks 221 and 219. The network element 227 terminates the tunnels 223, 225 and routes the traffic to the Internet 229.

Although BCP can be employed to transmit an Ethernet frame, it is an inefficient method. Both PPP and BCP must be supported on the network elements processing BCP data packets. In addition to the network elements, the CPE must also support BCP which is atypical. Moreover, BCP support is in addition to whatever tunneling protocol is to be used. Utilizing GRE to transmit Ethernet frames presents a more efficient and simple

method than BCP, but GRE does not have the robust signaling integral to L2TP.

Customers requesting certain information passed as attribute values pairs (AVPs) in L2TP lose the option of receiving such information with GRE.

SUMMARY OF THE INVENTION

The invention provides an apparatus and method for transmitting Ethernet data. According to one aspect of the invention, a method provides for transmitting Ethernet frames over a tunneling protocol.

In one embodiment of the invention, an Ethernet frame is received and transmitted over a non-homogenous tunnel, the tunnel having a plurality of sessions. Requested values are also transmitted over the non-homogenous tunnel.

In an alternative embodiment of the invention, Ethernet data included in an Ethernet frame is transmitted over a non-homogenous L2TP tunnel to a service provider. Upon receiving the Ethernet data, the service provider analyzes the Ethernet data.

In another embodiment of the invention, an Ethernet frame is encapsulated with L2TP, transmitted over a network, and decapsulated. The Ethernet frame is transmitted on one of a plurality of sessions in a non-homogenous tunnel running over the network. Attribute value pairs (AVPs) are transmitted over the session in relation to the encapsulated Ethernet frame.

These and other aspects of the present invention will be better described with reference to the Detailed Description of the Preferred Embodiment and the accompanying Figures.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention. In the drawings:

The invention may best be understood by referring to the following description and accompanying drawings that are used to illustrate embodiments of the invention. In the drawings:

Figure 1 (Prior Art) is a diagram of transmitting an Ethernet frame with Bridge Control Protocol (BCP) encapsulation.

Figure 2 (Prior Art) is a diagram of transmitting traffic from multiple subscribers who use different protocols with multiple Generic Routing Encapsulation (GRE) tunnels.

Figure 3 is a diagram of transmitting traffic from multiple subscribers who use different protocols over a single Ethernet capable Layer 2 Tunneling Protocol (L2TP) tunnel according to one embodiment of the invention.

Figure 4A is a diagram illustrating establishing an Ethernet capable tunnel between a LAC and LNS according to one embodiment of the invention.

Figure 4B is a diagram illustrating establishment of an Ethernet over L2TP session between a LAC and LNS according to one embodiment of the invention.

Figure 5 is a flowchart of a LAC processing traffic according to one embodiment of the invention.

Figure 6 is a flowchart for processing traffic received from a session in an Ethernet capable L2TP tunnel according to one embodiment of the invention.

DETAILED DESCRIPTION OF THE DRAWINGS

The present invention provides for a method and apparatus for transmitting Ethernet frames across a network. In the following description, numerous specific details are set forth to provide a thorough understanding of the invention. However, it is understood that the invention may be practiced without these specific details. In other instances, well-known protocols, structures and techniques have not been shown in detail in order not to obscure the invention.

endpoint uses less resources and can process more traffic than the same network element acting as a GRE tunnel endpoint in situations like that shown in Figure 3. Exemplary techniques for implementing non-homogenous L2TP tunnels that can carry Ethernet are described in more detail later herein..

Figures 4A-4B are diagrams illustrating establishment of an Ethernet capable tunnel and an Ethernet over L2TP session between a LAC and an LNS. Figure 4A is a diagram illustrating establishing an Ethernet capable tunnel between a LAC and LNS according to one embodiment of the invention. The LAC 401 creates and transmits an L2TP Start-Control Connection Request (SCCRQ) or tunnel request control message 405 to the LNS 403 for tunnel setup (it is understood that a network element can act as both a LAC and an LNS). The tunnel request control message 405 indicates to the LNS 403 that the LAC might transmit L2TP encapsulated Ethernet frames. The LNS 403 transmits Start-Control-Reply (SCCRP) or a tunnel reply control message 407 to the LAC 401 confirming its ability to process an Ethernet frame. If a tunnel is established between the LAC 401 and the LNS 403, then a session must be established to actually transmit data.

Figure 4B is a diagram illustrating establishment of an Ethernet over L2TP session between a LAC and LNS according to one embodiment of the invention. The LAC 401 transmits Incoming Call Request (ICRQ) or a session request control message 409 to the LNS 403. The request control message 409 indicates to the LNS that Ethernet frames are being transmitted in the L2TP session and indicates the Ethernet Media Access Control (MAC) Address of the LAC. When the LNS transmits traffic to the subscriber, the LNS uses this MAC address as the source address of the traffic being transmitted to the subscriber. After a session is established between the LAC and LNS, the LAC begins to transmit Ethernet over L2TP data packets to the LNS 403.

Figure 5 is a flowchart of a LAC processing traffic according to one embodiment of the invention. At block 500 a LAC listens for traffic on a bridged circuit or Ethernet port. At block 501, the LAC receives traffic on the bridged circuit or Ethernet port. The

LAC determines if the received traffic is to be tunneled at block 503. The LAC then queries a database such as RADIUS for tunnel parameters. If a failing event occurs which prevents the LAC from receiving tunnel parameters, then at block 507 the bridged circuit or Ethernet port is paused for a configured amount of time before returning control to block 500. A failing event can include a communication failure between the LAC and the database, the database not finding tunnel parameters, etc. If the LAC does attain tunnel parameters, then at block 509 the LAC attempts to establish a tunnel with an LNS. If the tunnel setup fails, then control goes to block 507. If the tunnel is established, then the LAC requests a session over the tunnel with the LNS and indicates Ethernet and the LAC's MAC address at block 511. If the session setup fails, then control flows to block 507. If session setup is successful, then the LAC begins to transmit the data traffic to the LNS at block 513. Although the described embodiment pauses after all three failing scenarios, alternative embodiments can pause for select failing events or not pause.

Figure 6 is a flowchart for an LNS processing traffic received over a session in an Ethernet capable L2TP tunnel from a LAC according to one embodiment of the invention. At block 601, an LNS receives an Incoming-Call-Request (ICRQ) or session control message from a LAC to establish an L2TP session. An AVP in the ICRQ control message indicates the session type to be transmitted from the LAC. At block 603, the LNS determines if the session request control message indicates Ethernet. If the control message does indicate Ethernet, then a virtual circuit structure is created indicating Ethernet encapsulation at block 607. Otherwise, the LNS creates a virtual circuit structure at block 605 indicating a different encapsulation. At block 608 the LNS receives an L2TP data packet over the session. When the LNS receives the L2TP packet from the session, a decapsulation routine removes the L2TP header from the packet to get a payload at block 609. The LNS then associates the decapsulated payload with the virtual circuit structure at block 611. The LNS processes the payload as indicated by the virtual circuit structure at block 613. If the virtual circuit indicates Ethernet

encapsulation, the virtual circuit structure creates the impression to the system of the LNS that the payload is a data packet received on an Ethernet port. If the virtual circuit structure indicates a different encapsulation, then it creates the impression to the LNS system that the payload was received on a circuit configured for the different encapsulation.

The capability to transmit Ethernet frames over L2TP tunnels increases functionality of a network element supporting this extension to L2TP. This extension of L2TP covers the two network layer protocols most likely employed by subscribers. The wholesale network provider can employ L2TP to provide ISP customers individual data streams for each subscriber accessing the network through the wholesale network provider's network element. L2TP also enables a network provider with robust signaling functionality. With L2TP, the network provider can satisfy a customer request for certain information to be passed as AVPs. The flexibility of L2TP allows the network provider to add and remove AVPs in response to the changing needs of customers. Other protocols such as GRE do not support this signaling functionality. Moreover, additional protocols are unnecessary as with BCP.

In addition, GRE is limited to IP media which is more complex and has more overhead than other media such as Frame Relay. In contrast, L2TP can be carried over any media.

The techniques shown in the figures can be implemented using code and data stored and executed on computers. Such computers store and communicate (internally and with other computers over a network) code and data using machine-readable media, such as magnetic disks; optical disks; random access memory; read only memory; flash memory devices; electrical, optical, acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.); etc. Of course, one or more parts of the invention may be implemented using any combination of software, firmware, and/or hardware.

While the invention has been described in terms of several embodiments, those skilled in the art will recognize that the invention is not limited to the embodiments described. In an alternative embodiment, the data stored in an Ethernet header can be transmitted instead of the Ethernet frame. The Ethernet payload could be decapsulated from the Ethernet frame. The data in the Ethernet frame could be stored as values in the L2TP header. Hence, the Ethernet information could be transmitted with L2TP instead of encapsulating the Ethernet frame with L2TP. In another embodiment of the invention, Ethernet could be transmitted over L2TP and PPP. A virtual PPP client would run on the LAC. The Ethernet frame could be encapsulated with the virtual PPP client, and the Ethernet frame transmitted within PPP encapsulation without extending L2TP.

The method and apparatus of the invention can be practiced with modification and alteration within the spirit and scope of the appended claims. The description is thus to be regarded as illustrative instead of limiting on the invention